

What is Claimed:

1. A method for performing durable reads of a data page in a system that permits lazy commit transactions, said method comprising flushing a transaction log associated with said data page prior to said data page being read as a part of a durable read operation.
2. The method of claim 1 further comprising:
marking said data page when modified by a lazy commit transaction; and
unmarking said data page when said commit log associated with said data page is flushed.
3. The method of claim 2 wherein the step of flushing a transaction log associated with said data page prior to said data page being read by a durable read operation occurs when said data page is marked, and wherein said method further comprises reading an unmarked data page as part of a durable read operation without first flushing said transaction log associated with said data page.
4. The method of claim 3 wherein the step of marking a data page when modified by a lazy commit transaction comprises writing a value of a bit associated with said data page.
5. The method of claim 4 wherein the bit is stored in said data page.
6. The method of claim 4 wherein the bit is stored in a reference table.

7. The method of claim 3 wherein the step of marking a data page when modified by a lazy commit transaction comprises recording, in a reference location associated with said data page, a copy of a log sequence number from said transaction log and corresponding to said modification of said data page by said lazy commit transaction.
8. The method of claim 7 wherein said copy of the log sequence number is stored in said data page.
9. The method of claim 7 wherein said copy of the log sequence number is stored in a lazy commit reference table.
10. The method of claim 7 wherein the copy of the log sequence number is used to identify the lazy commit transaction in order to cause said lazy commit transaction to effect the step of flushing said transaction log associated with said data page and unmarking said data page when said data page is flushed.
11. A computer-readable medium having computer-readable instructions for performing durable reads of a data page in a system that permits lazy commit transactions, said computer-readable instructions comprising instruction for flushing a transaction log associated with said data page prior to said data page being read by a durable read operation.
12. The computer-readable medium of claim 11 further comprising instructions for:

marking said data page when modified by a lazy commit transaction; and
unmarking said data page when said transaction log associated with said data page is
flushed.

13. The computer-readable medium of claim 12 wherein the step of flushing a transaction log associated with said data page prior to said data page being read by a durable read operation occurs when said data page is marked, and wherein a durable read operation can read an unmarked data page without first flushing said transaction log associated with said data page.
14. The computer-readable medium of claim 13 wherein the instructions for the step of marking a data page when modified by a lazy commit transaction further comprises instructions for changing a value of a lazy commit bit associated with said data page.
15. The computer-readable medium of claim 14 further comprising instructions for the lazy commit bit to be stored in said data page.
16. The computer-readable medium of claim 14 further comprising instructions for the lazy commit bit to be stored in a lazy commit reference table.
17. The computer-readable medium of claim 13 wherein the instructions for marking a data page when modified by a lazy commit transaction further comprises instructions for recording a copy of a log sequence number, from said transaction log and corresponding to said modification

of said data page by said lazy commit transaction, in a reference location associated with said data page.

18. The computer-readable medium of claim 17 further comprising instructions for said copy of the log sequence number to be stored in said data page.

19. The computer-readable medium of claim 17 further comprising instructions for said copy of the log sequence number to be stored in a lazy commit reference table.

20. The computer-readable medium of claim 17 further comprising instructions for the copy of the log sequence number to be used to identify the lazy commit transaction in order to cause said lazy commit transaction to effect the step of flushing said transaction log associated with said data page and unmarking said data page when said data page is flushed, as well as instructions for said lazy commit transaction to flush said transaction log associated with said data page and unmark said data page when said data page is flushed.

21. A system for performing both lazy commit transactions and durable reads, said system comprising:

- a plurality of data pages;
- a plurality of transaction logs associated with each of said plurality of data pages; and
- a durability subsystem that flushes a transaction log, from among said plurality of transaction logs, associated with a data page, from among said plurality of data pages, prior to said data page being read by a durable read operation.

22. The system of claim 21 further comprising a lazy commit subsystem whereby lazy commit transactions mark said data page when said lazy commit transaction modifies said data page.

23. The system of claim 22 further comprising a durable read subsystem whereby said durable read operation, when executing the process of reading said data page, checks whether said data page has been marked and, (a) if so, flushes a transaction log associated with said data page, unmarks said data page, and reads a set of data from said data page, and, (b) if not, reads the set of data from said data page without first flushing said transaction log associated with said data page.

24. The system of claim 23 wherein the data page comprises a lazy commit bit associated with said data page that is changed when said data page is modified by said lazy commit transaction.

25. The system of claim 24 wherein the lazy commit bit is stored in said data page.

26. The system of claim 24 wherein the lazy commit bit is stored in a lazy commit reference table.

27. The system of claim 23 further comprising a marking subsystem which records a copy of a log sequence number, from said transaction log and corresponding to said modification of said

data page by said lazy commit transaction, in a reference location associated with said data page when said data page is marked.

28. The system of claim 27 wherein marking subsystem uses the copy of the log sequence number to identify the lazy commit transaction in order to cause said lazy commit transaction to effect the step of flushing said transaction log associated with said data page and unmarking said data page when said data page is flushed.